

Applications and design considerations for optical interconnects in VLSI

L. Bergman, A. Johnston, and R. Nixon
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109

S. Esener, C. Guest, P. Yu, T. Drabik, M. Feldman, and S. H. Lee
Department of Electrical Engineering and Computer Sciences
University of California, San Diego
La Jolla, CA 92093

Abstract

This paper introduces new applications and design tradeoffs anticipated for free space optical interconnections of VLSI chips. New implementations of VLSI functions are described that use the capability of making optical inputs at any point on a chip, and take advantage of greater flexibility in on-chip signal routing. These include N-port addressable memories, CPU clock phase distribution, hardware multipliers, dynamic memory refresh, as well as enhanced testability. Fault tolerance and production yields may be improved by reprogramming the optical imaging system to circumvent defective elements. These attributes, as well as those related to performance alone, will affect the design methodology of future VLSI ICs. This paper will focus on identifying the design issues, their possible solution, and their impact on VLSI design techniques.

Introduction

The purpose of this paper is to outline our initial thoughts on how free space optical interconnections might be effectively used in existing microcircuits. Potential circuit applications are mentioned, and design considerations for the main components are discussed as well as some test chip layouts incorporating optical input gates. At this stage, few experimental results are available; the paper emphasizes current study efforts which will lead to experiments to be reported on at a future date.

Following this, the general background and types of applications of optical interconnects are first described, concentrating on systems requiring only inputs to silicon circuits from off-chip light sources because of present technology limitations. A few specific applications involving microcircuits in current use are briefly described. Applications to chip level microcircuits and to wafer scale integration are discussed.

In the following section, detailed design considerations for the three main components of an optical interconnect system, that is, the light sources, the holographic optical element (HOE), and the detectors and associated circuits forming an optically addressed gate are discussed. Some early experiments with HOE generation are summarized. In the case of the optical input gates, detector designs based on standard silicon MOS fabrication processes in use for custom microcircuit fabrication are outlined. A test chip incorporating these detectors was fabricated, and is described.

In the final section, our conclusions and ideas for future work are briefly outlined.

Background

Over the past two decades, tremendous strides have been made in microelectronic design and fabrication techniques, that have resulted in increased speed and component density for VLSI chips. Ultimately, it is expected that the maximum number of components that can be fabricated per unit area will be constrained by a physical limit of about 0.1- μ m feature sizes. Long before this occurs, however, difficulties are expected to arise in (1) testing such complex ICs, (2) providing enough off-chip interconnects for large ICs, and (3) overcoming handlimiting effects resulting from driving interconnection lines both on and off chip.¹⁻³ The trend toward faster GaAs materials technology is expected to exacerbate these problems still further.⁴ In addition, even without extrapolating to future device densities, speed and layout constraints are important factors in present VLSI design.

Optical chip interconnections have been suggested by Goodman as one means of alleviating these problems and have been discussed in depth by him and his coauthors in two recent papers.^{5,6} Unlike electrical connections, optical interconnects are insensitive to mutual

interference effects, free from capacitive loading and planar (or quasi-planar) constraints, and can be reprogrammable. Both waveguide optics and free-space optical interconnections have been proposed. However, the former has handicaps that severely limit its use, in particular, high attenuation, high bending loss, and constraint to a plane if fabricated in an integrated optics form.

Although alignment problems are more critical if sources and detectors cannot be directly incorporated into the microelectronic fabrication process, free space optical interconnects suffer fewer of these disadvantages and further offer the possibility of a broadcast mode (one to many) connection.

As shown in Figure 1, a holographic optical element (HOE) optically connects a given source-detector pair existing on either the same chip or two different chips. Many applications exist, and a few specific examples are described in greater detail below.

In attempting to implement optical interconnection of electronic circuits, the present limitations of both technologies must be taken into account. First, no method of fabricating useful optical sources in silicon is currently known, though optical detectors in silicon are commonplace. Second, integration of complex circuitry in GaAs is just beginning to occur. Therefore, a reasonable initial approach is to use off-chip GaAs sources to provide inputs to detectors, and circuitry integrated in silicon. Eventually, the sources must be brought onto the electronic chip either through fabrication of the electronics on GaAs, or by hybrid bonding of GaAs sources onto Si substrates. More general interconnection arrangements will then be possible. Nevertheless, silicon microcircuits can provide a convenient test vehicle for evaluating optical interconnects and their impact on IC layout and circuit architecture.

Two general types of applications can be imagined, even if the constraint of using off-chip sources is observed. Broadcasting of one signal to many locations on a chip, such as clock distribution, is one potential application type. Implementing the row and column address lines, as well as strobe lines, on a large RAM chip is a similar example that allows considerable area previously used for interconnections to be eliminated from the chip. It can also reduce or eliminate the delays inherent in resistive polysilicon transmission lines. A second application type is any chip whose functionality is constrained by limitations on the number of input/output pins. By providing optical inputs to chips such as large binary multipliers, or Programmable Array Logic devices (PALs), all available pinout positions at the periphery of the chip can be used for outputs. An example combining both advantages above is to use optical inputs to provide a parallel load capability to a systolic array processor. All elements of the array could be provided in parallel with programming information or initial data. In addition, optical interconnects would relax the layout of the array as nearest neighbor constraints need not be observed.

The speed of present VLSI circuits is limited to a large extent by the interconnections between the input pads and the data processing area. The use of optical input can improve the VLSI circuit speed, as the speed of optical sources and receivers can be comparable in performance (100 ps to 1 ns rise time) with the other fast elements on the chip. Circuits such as optical receiver modules are RC-bandlimited by stray capacitance and stray resistance, and can also be further optimized in an on-chip form. Speed constraints imposed by interconnection lines will be even more severe in ultrafast GaAs logic, and as a result, optical techniques may be much more important there in the future.

Specific applications

The following examples are presented in order to illustrate applications of optical interconnects in specific circuits, most made using current state-of-the-art VLSI silicon foundry technology. These applications include both chip scale and wafer scale integration.

Chip scale

N-port addressable memory. One of the most interesting applications of optical interconnects is in constructing new functional elements within current silicon VLSI technology that would otherwise be impossible to implement due to limited package connections. One such case is an n-port memory. This is a RAM memory system that allows n read and/or write operations to be made simultaneously. It is typically used in a multiprocessor environment where many CPUs and peripheral units must access the same primary memory. Sharing memory resources in this manner enhances utilization in large systems and provides increased inter-processor communications efficiency.

The principal reason these ICs have not been fabricated previously is due to the limited number of address and data lines that can be accommodated in a typical IC package. With optical interconnects, this number can be increased by perhaps one order of magnitude or

more (>10 ports). These ICs will also require additional logic for contention resolution of the data channels.

CPU clock phase distribution. A typical microprocessor VLSI chip contains a great deal of synchronous logic and requires the distribution of a multiphase clock throughout most of its architecture. Electrical routing of these signals uses silicon real estate inefficiently and worse, introduces unknown propagation delays and crosstalk in the clock distribution network. A number of design and test chip iterations are usually required to remove these bugs and to improve the speed. Optical interconnects would provide the ability to transmit any clock phase to any location on the chip with less phase skew and crosstalk. The net result would be a more compact design that operates faster and performs closer to the intended design model. Also, production yields may improve because of the reduced crosstalk and tighter tolerances that can be achieved.

Hardware multiplier. The arithmetic multiply operation is one of the most frequently used operations in image processing, engineering simulation, and numerical analysis. Since it also requires much more time to perform than addition or other boolean operations, main-frame computer manufacturers have gone to great lengths to build special hardware to reduce multiplication time as much as possible. Even in moderately low-speed MOS/TTL computers (<20MHz), faster semiconductor technologies such as Schottky TTL or ECL have been employed to reduce the multiply execution time to the magnitude of the addition time (typically a factor of 4 to 10 speedup). For example, a 16 x 16 bit multiply operation typically takes about 150 ns in MOS, but only 25 ns in 100K ECL and 10 ns in GaAs logic.

The next generation of supercomputers will possess clock rates of several GHz, and will likely perform addition operations in 1-4 ns. This suggests that a hardware multiply function of equal speed should be provided. Unless still faster semiconductor technologies become available, designers will be forced to examine alternative architectures for implementing this function. One promising candidate is to revise the multiply algorithm so that parallelism is exploited more. For example, a look-up table approach would permit a result to be obtained in essentially one memory cycle time at the expense of using more silicon real estate and more interconnects. Optical interconnects would be essential in accommodating the large number of multiplicand and product line connections to memory cells and minimizing their crossings. However, finite ROM sizes would restrict this technique to relatively small numbers of multiplicand and product lines. Another approach that exploits parallelism would be to use optical interconnects perhaps on a wafer scale to provide the multiplicity of data line connections required to calculate large integers from arrays of conventional IC multiplier elements.

Dynamic RAMs & shift registers. Dynamic memories require the distribution of many different clock phases for addressing, refresh, and sensing operations. Optical interconnects could be used to reduce clock delay and skew in dynamic RAMs for an overall speed gain. Dynamic shift registers store information by circulating bits in a bucket-brigade fashion. Each cell must be clocked periodically to both retain and shift its information to the next cell. Typically, the clock signal is the same for all cells thus making a broadcast clock distribution an efficient method of providing timing for the chip. This is ideally suited for an optical interconnection scheme. Using current silicon VLSI technology, a single laser source could broadcast a clock signal via an HOE to a two-dimensional array of silicon detectors, one positioned at each memory cell. By eliminating long strobe lines on the chip, higher clock speeds would be possible.

Future enhancements. Microcircuit testing techniques based on optical access to specific locations on a chip can be developed. Further, in a system using optical interconnects, a partially defective chip could be salvaged by customizing the holographic element used with it to bypass failed circuits and instead direct the optical inputs to redundant back-up elements. An extension of this capability is to optically reconfigure inputs in real time to compensate for failed parts or to select among several possible functions.

The ultimate goal of reconfigurable interconnects is to be able to change the interconnect matrix as quickly and as freely as needed. Such a capability will allow optical interconnection to improve on many of the functions presently implemented on a limited scale with electronics, such as routing data between processing elements based on data-dependent decisions, and multiplexing and demultiplexing information. The potential applications for reconfigurable holographic optical elements are so compelling that investigation of alternatives for programmable interconnections will be of great interest in the future.

Wafer scale

Integration at the wafer level is the ultimate extent to which monolithic device fabrication can be carried, since devices are not cut apart at all. Both the advantages and disadvantages of higher level integration are greatly pronounced. High circuit density and

the removal of restrictions on the partitioning of a system into chip-sized blocks make the performance potential of wafer-scale integration (WSI) very great indeed. However, problems of yield, the difficulty of clock distribution for synchronous systems, and the real-estate and speed penalties paid for global intra-wafer connections detract significantly from the benefits of WSI and impose constraints on the architectures of systems suitable for WSI realization.

The high cost of global interconnection has led to the design of planar architectures which require only local data movement, such as tree or mesh-connected arrays of processing elements. The global timing problem can be eliminated at the expense of adopting asynchronous communication protocols and asynchronous architectures in general. The question of yield is very serious, and attempts to accommodate low wafer yield center around fault tolerant array architectures consisting of processing element arrays with connection paths that are custom-configurable to some extent, for example by means of laser trimming. These approaches rely on redundancy and may not make full use of all functional processing elements on the wafer.

The application of optical interconnects can remove most of these constraints on array architectures suitable for WSI. Optical clock distribution enables global synchronism. Global data movement can be accomplished optically, and as mentioned above, the interconnection hologram can be fabricated to connect functional processing elements on a given imperfect wafer, thereby accommodating the effects of poor yield.

Some new constraints on interconnection are imposed by practical considerations. The density of optical sources on the wafer is limited by thermal dissipation density, real estate issues, and practical hologram requirements. For a given set of fabrication technologies, there will be a break-even distance Δ over which an electrical and an optical connection would be equally expensive. Optical interconnections should be used over distances greater than Δ , and metal lines should be used otherwise, as a rule of thumb. This suggests that the size of an array element should be on the order of Δ .

A practical WSI processor might have hundreds of processing elements, each with an area of 10 mm², and with several optical inputs and one or two optical outputs. Many highly parallel high-wire area architectures requiring global data movement, such as hypercubes and butterfly machines, as well as planar architectures, such as tree machines and mesh-connected arrays, fit well into such a context.

Detailed design considerations

Holographic optical element

The holographic optical element (HOE) defines and implements the interconnection pattern between sources and detectors. It determines the fan-in, fan-out, and distribution of signals. Use of an HOE rather than electrical connections or optical fibers allows a large amount of freedom in designing the interconnection pattern and in reprogramming the pattern if necessary.

There are specific requirements to be met for an HOE used for interconnection of VLSI circuits. Light from the laser diode sources will reach the hologram as ellipsoidal expanding wavefronts. The direction and position of light arrival at the hologram will be strongly influenced by the geometrical arrangement of sources. Each signal wavefront must be reflected, focussed and directed to one or more detector gates. The pattern of distribution for each signal may be different. The HOE must be highly efficient to ensure that most of the light from each source is used to provide signal to the detectors. High HOE efficiency reduces source power consumption and allows better noise tolerances for the detectors. Optical noise in the interconnections will arise from two aspects of the HOE. First, to whatever extent the HOE is not perfectly efficient, the undiffracted light produces a diffuse unfocussed background light level at the detectors. Second, all holograms produce some amount of stray light, resulting from stochastic effects in the recording process and, in the case of computer generated holograms, quantization effects in the encoding. Optical noise not only constrains signal level margins, but can lead to signal crosstalk as well. The hologram must produce tightly focussed spots at the detector locations. To maximize the signal reaching the detectors and to minimize optical effects on surrounding circuitry, the spot size should closely match the size of the detectors, i.e., about 10 by 10 microns. One additional consideration in the design of the hologram is facilitation and maintenance of alignment. To allow automated assembly techniques, the hologram must provide a means of being aligned with the sources and detectors in another plane; and it must be tolerant of slight misalignment due to imprecise assembly and thermal distortions. Furthermore, some distribution of wavelengths from different source lasers must be tolerated, as well as wavelength wander in each source over time.

The desirability of having sources and detectors in a single plane due to packaging

requirements and minimization of electrical pathlengths dictates that a reflective hologram be used. The hologram will lie in a plane parallel to the source and detector plane at a distance of 5 to 10 mm. A true reflection hologram requires fringes that run generally parallel to the surface of the hologram. While technology for computer generation of such structures is under development, an alternate approach of combining a mirror and a surface relief hologram into a single element will be used. This structure can take the form of a surface relief pattern etched into a silicon surface, which is naturally reflective, or a reflective metal may be coated onto less reflective patterned substrates.

Considerable flexibility is inherent in the use of free space optical interconnects. Each signal source can be connected to several detectors. The connection pattern for each source can be quite different. Each detector may also receive inputs from several signal sources. A simple Boolean OR of the received signals is performed by the detector. Sources and detectors will be rather arbitrarily positioned on their respective chips. To provide the required interconnection flexibility, computer generation of the HOE pattern is called for. All interconnection paths are combined into one computed pattern that is then recorded by automated means. This approach has two further advantages over natural holographic elements. Source and detector position information can be automatically extracted from commonly used CAD tools to provide input specification for the HOE pattern computation. Also, the computed holographic pattern can be fabricated using methods compatible with existing integrated circuit manufacturing facilities.

The approach to fabricating practical HOEs for VLSI interconnect involves a series of steps, making use at each step of the experience gained previously. Initial experiments were performed using thin absorption transmission holograms printed on photographic film using a laser scanner system. These have advantages of immediate availability and short fabrication time for prototyping new versions. These holograms are formed from overlapped binary Fresnel zone plates. Sample holograms have been produced to provide connection of one source to one detector and one source to five detectors. These are available for preliminary experiments for integrating sources, holograms and detectors into a working system. Diffraction efficiencies of five to six per cent have been achieved.

The photographic film hologram was then contact printed onto a layer of photoresist on a silicon substrate. The photoresist is developed and the silicon etched by chemical means. After removal of the remaining photoresist, a reflective surface relief hologram is present in the surface of the silicon. Because this is a phase hologram, much higher diffraction efficiencies (theoretical maximum is 40%) can be expected.¹⁰ Experiments are in progress to evaluate these holograms.

In the future we plan to use electron beam lithography to write the hologram pattern. This will produce holograms with much higher space-bandwidth product (SBP) than can be produced by the laser scanner system. The practical effects of increased SBP are more complex interconnection patterns and lower f-number HOEs leading to more compact system configurations. Further development paths that will be investigated are high aspect ratio surface relief holograms that can provide diffraction efficiencies approaching 100%, and computer generated volume holographic optical elements to allow angular multiplexing of connection patterns.

Laser sources

Light sources are required to convert electrical signals to optical form. Due to the coherence requirements of HOEs, semiconductor laser diodes should be used for illuminating the hologram. The compact size and low electric power requirements of these laser diodes will ease the problems associated with the assembly and integration of the hologram and the VLSI chip. One significant advantage of the laser diode is its potential for high modulation frequency—as high as 18 GHz has been demonstrated. However, for the success of this program, several key requirements of the laser diodes must be considered. First of all, since several independent light sources are needed as inputs for a VLSI chip, an array of independently addressable laser diodes is required. Each of these lasers is DC biased to just below the threshold voltage, and the input electrical signal is superimposed on this bias to trigger the laser. For both versatility and ease of assembly, the output beams from the laser diodes must be directed out of the plane of the mounting substrate. This poses some difficulty for assembly, as the laser diode has typical dimensions of 200 μm (length) x 200 μm (width) x 75 μm (thickness), and the laser cavity is along the length dimension. To alleviate this problem, the laser diode is first mounted on another substrate and this substrate is then mounted vertically on the silicon in such a manner that the laser beam is directed out of the plane. Since it is desired to address several input gates on the VLSI chip from the output of each laser, high efficiency laser diodes are needed. A typical required peak light output from the laser of 5-10 mW is estimated, based on the expected hologram efficiency, the silicon detector dimensions and output circuit requirements.

The most critical requirements on the laser diode are the spectral width and the wavelength stability of the laser beam. In order to maintain alignment tolerances, it is essential to match the holographic grating periodicity as closely as possible to the laser wavelength. To meet this end, an index-guided laser with built-in structures for wavelength selection and stabilization is considered. In addition, heat sinking with feedback to the drive circuitry should be provided to fine tune the emission wavelength. Different approaches for fabricating an array of lasers will be considered. One can either mount separate laser diodes on a substrate or, alternatively, mount a monolithic integrated array of lasers. With an array of discrete lasers, one can control individual laser emission separately with separate heat sinks and associated circuitry. Also, it is technologically easier to fabricate single laser diodes than an array of laser diodes. However, with monolithic arrays, one has better control over the laser spacing and orientation.

The light sources that will be used in our experiments will be GaAs/GaAlAs lasers emitting in the wavelength range from 820 to 860 nm. The laser chosen is a buried heterostructure window laser with a typical active region cross section of $2\text{ }\mu\text{m} \times 0.15\text{ }\mu\text{m}$. This gives, in accordance with the theory of an index guided waveguide, an angular divergence of the output beam of $25^\circ \times 35^\circ$. In addition, this laser possesses both a stable high power output and good modulation characteristics. The typical threshold current of the laser is 20-30 mA and the power consumption when on is about 0.2 watt.

Detectors & detection circuits

The final operation in optical interconnection is the detection of signals generated by optical sources and distributed by the HOE to appropriate locations on the silicon chip. In this section, the design trade offs for the detectors are outlined, and then the design and fabrication of a set of CMOS detector test circuits are discussed in more detail.

For optical interconnects to be viable, the detectors must meet certain requirements. First, the process used to fabricate the detector must be compatible with the processes used in fabrication of the electronic circuits on the VLSI chip. Extra process steps will increase the chip cost and decrease the yield. Second, the detector circuit area must be small enough to not significantly increase the chip size or restrict the layout. Large chip size may decrease yield and speed (longer local interconnects), while drastic changes in layout will increase the development time of the VLSI chip. Third, the detection area must be designed such that spatial drifts of the beam due to the frequency shift of the optical sources can be tolerated. However, the detector should also be large enough to avoid problems associated with noise limitations. The noise performance of the complete detector circuitry should maintain a specified error rate. Fourth, the detectors should be sensitive enough to the incoming radiation wavelength to generate operating voltages compatible with existing logic families (e.g., CMOS). In some cases, where the photocurrent is too small to charge the load capacitance at the required rate, or when the standard logic levels cannot be reached, a buffer preamplifier circuit can be incorporated. Finally, the detection circuitry must be fast enough not to limit the optoelectronic VLSI system speed.

Since a RAM chip previously developed at JPL¹¹ will be used as a first example to focus our study of optical interconnection techniques, we shall investigate detector structures compatible with the CMOS process with which it is fabricated. The RAM chip was fabricated using a₁₂ 3-micron CMOS-bulk p-well process, through the MOS Implementation Service (MOSIS).

MOSIS is a fast turnaround "silicon broker" that serves the requirements of DARPA's VLSI research program and involves participants from several universities and other research and development organizations. DARPA has established the MOSIS system at USC/ISI, located at Marina Del Rey, CA. The service provides the interface between the designer and a number of fabrication lines. The cost per project is kept at a reasonable level by merging many die types onto one wafer. MOSIS currently supports 2-micron NMOS, 3-micron CMOS-bulk, and 3-micron CMOS-SOS as standard technologies. MOSIS plans to support smaller feature sizes, as well as GaAs in the near future.

Three different detector types, shown in Figure 2, $p^+n^-p^+$ lateral phototransistors, lateral $p^+n^-n^+$ photodiodes and vertical p^+n^- photodiodes are compatible with this process without requiring any additional process steps. Lateral $p^+n^-p^+$ phototransistors can provide larger photocurrents due to their internal photogain but will have slower response time. Lateral $p^+n^-p^+$ photodiodes will provide high speed because of small associated capacitances but will have poor sensitivity. In a MOSIS CMOS process the junction depth of the p^+ and n^- layer is only 0.4 μm , while laser radiation at 820 nm requires about a 10 μm -thick Si layer for absorption. The responsivity of these lateral diodes can be improved by using the p-well diffusion process step of the MOSIS process to achieve a thicker photoactive region. However, these lateral $p^+n^-p^+$ photodiodes will have a poor ratio of the effective detection area to device area due to their interdigitated layout. Therefore their expected sensitivity remains quite low. The vertical p^+n^- photodiode with a thin p^+

layer will provide good sensitivity as well as acceptable response time for many applications and may be best suited for medium-speed CMOS applications. The speed of the vertical p-n photodiode will be limited by its large parasitic capacitance and diffusion effects.

The nature of the optical interconnection application will dictate the requirements on the detector circuit since the charging or discharging speed depends on the effective load capacitance it drives. The load capacitance consists of the detector shunt capacitance and the input capacitance of the driven CMOS gate. In summary, if the input capacitance of the CMOS gate is large, a p-n photodiode becomes the detector of choice. If the loading CMOS gate is small, with a low input capacitance, the p⁺pn⁻ photodiodes should provide higher speed operation. When the photocurrent is not sufficient to charge and discharge the load capacitance, a buffer CMOS gate should be used. The detection circuit then consists of a photodiode, its load, and a CMOS buffer gate as shown in Figure 3. Because of the standard biasing of bulk CMOS circuits (n⁻ being positive), it is advantageous to connect one side of the detector to the positive rail. The n⁻ bulk silicon is then used for detection. The cross-section of a CMOS detection circuit implemented with a vertical p-n photodiode, an enhancement mode PMOS load, and a CMOS buffer gate is shown in Figure 4. The enhancement mode PMOS load reduces the required voltage swing by setting the low output voltage state of the detector to the PMOS threshold value.

Detection circuits using vertical p⁺n⁻ and p⁺pn⁻ photodiodes with CMOS buffer gates have been designed and fabricated using the MOSIS process. The diagram shown in Figure 5 illustrates the layout of the optical input gate circuit shown schematically in Figure 3. The illustrated detector is the vertical p-n photodiode. These circuits are incorporated in a test chip, shown in Figure 6, which is intended for use in experiments to determine the performance of the optical interconnect components. These chips are currently being evaluated, and will also be used in testing of the HOEs.

As shown in Figure 6, twelve optical input gates in a linear array appear at each of five sites on the chip. Each of the twelve detectors has a set of different design parameters and can be selected independently from others. The five arrays are placed at the corners of a 3 mm square and at its center.

The design of the optical gates was optimized for an incident light power of .1 mW at 820 nm and an optical beam diameter of 15 μm. The detector area was 9 x 9 μm² and a 6000 Å thick field oxide was left on the detector area to serve as an antireflection coating. The photocurrent produced by the photodiode can be estimated by

$$I_{ph} = \frac{P_{inc} A_d}{A_b} \frac{q(1-r)}{h\nu} (1 - e^{-\alpha d}) \quad (1)$$

where P_{inc} is the optical beam power, A_b is the beam area, A_d is the detection area, q is the electronic charge, h is Planck's constant, ν is the optical frequency, and d is the absorption layer thickness. The reflection and absorption coefficients are given by r and α respectively. Under the above specified conditions, a 20 μA photocurrent can be expected to flow.

Under the MOSIS process specifications and for the above geometries, the parasitic shunt capacitance C_d of the p-n photodiode is about 16 fF. The input capacitance C_{in} of a minimum geometry CMOS gate is about 8 fF and an input voltage swing V_{pp} of 2.25 V centered at 2.5 V must be applied to this gate to produce acceptable output levels. The charging time t_c of the detection circuit with a minimum size CMOS gate can then be estimated to be

$$t_c = \frac{V_{pp}(C_d + C_{in})}{I_{ph}} = 2.5 \text{ ns.} \quad (2)$$

Higher speeds should be achieved with p⁺pn⁻ photodiodes with smaller parasitic shunt capacitance. The discharging time can be tailored through the optimization of the PMOS transistor.

Conclusions

We have described in general terms a number of potential applications for optical interconnections in VLSI, more from the perspective of a VLSI circuit designer than that of an optoelectronic components expert. This leads into a qualitative discussion of the design of the three key components, the HOE, the light source, and the detector or optical input gate. Fabrication of test circuits using CMOS technology will allow experimental evaluation of optical performance of a realistic HOE, and measurement of stray light levels. The

power budget and efficiency will also be measured, as well as detector noise limits and response time.

These results will be reported on in future publications. Hopefully, they will contribute toward understanding how to effectively apply optical interconnection techniques in advanced microcircuits, and in understanding what circuits or what types of circuits can benefit most from this technology.

Acknowledgments

The authors would like to thank Dr. S.T. Eng for a number of helpful discussions dealing with both free space and fiber media in the early stages of formulating this work on optical interconnects for VLSI. Many of his ideas contributed to the effort, in both the technical and programmatic areas. We would also like to thank H. Sayah for laying out the test circuits and interfacing with the MOSIS system, C. Pina for providing valuable support and advice on selecting candidate test chips, and Dr. W. Wu for her assistance in preparing the manuscript.

This work was performed in part at the Jet Propulsion Laboratory and in part at the University of California, San Diego, under a contract from the Defense Sciences Office at DARPA.

References

1. P. S. Ho, "Some basic problems in VLSI interconnect," Third Int'l Symposium on Very Large Scale Integration Science and Technology, eds. W. M. Bullis and S. Broydo, Proceedings vol. 85-5, the Electrochemical Society, Pennington, NJ, May 1985, p. 146.
2. Anis Husain, "Optical Interconnect of Digital Integrated Circuits and Systems," SPIE Proceedings, Vol. 466-02, Jan. 26, 1984.
3. A. S. Bergendahl, J. F. McDonald, R. H. Steinworth, and G. F. Taylor, "Thick Film Stripline Micro Transmission Lines for Wafer Scale Integration," Third Int'l Symposium on Very Large Scale Integration Science and Technology, eds. W. M. Bullis and S. Broydo, Proceedings Vol. 85-5, The Electrochemical Society, Pennington, NJ, May 1985, p. 175.
4. B. K. Gilbert and D. J. Schwab, "Will present design approaches fulfill performance requirements for high clock rate digital signal processors based on GaAs?," SPIE Proceedings 466-01, Jan. 26, 1984.
5. J. W. Goodman, F. I. Leonberger, S. Y. Kung, and R. A. Athale, "Optical Interconnections for VLSI Systems," Proceedings of IEEE, 72, 850 (1984).
6. Raymond K. Kostuk, Joseph W. Goodman, and Lambertus Hesselink, "Optical Imaging applied to microelectronic chip-to-chip interconnections," Appl. Opt. 24, 2851 (1985).
7. Ravinder K. Jain, "Switching Characteristics of Logic Gates Addressed by Picosecond Light Pulses," IEEE Journal Quantum Electr., QE-19, 658 (1983).
8. F. J. Henley, "Functional Testing and Failure Analysis of VLSI Using a Laser Probe," Proceedings of Custom Integrated Circuits Conference, May, 1984.
9. S. H. Lee and J. Leger, "Signal Processing Using Hybrid System," Ch. 4 in Application of Optical Fourier Transforms, Ed. H. Stark, Academic Press (1982).
10. N. C. Gallagher, Jr., J. C. Augus, F. E. Caffield, R. V. Edwards, J. A. Mann, Jr., "Binary Phase Digital Reflection Holograms: Fabrication and Potential Applications," Applied Optics, Vol. 16, p. 413, 1977.
11. M. G. Buehler, B. T. Moore, and R. H. Nixon, "Parameter Extraction from Spaceborne MOSFETS," IEEE Trans. on Nucl. Sci., December, 1985.
12. "The MOSIS System (what it is and how to use it)," Report No. ISI/TM-84-128, Information Sciences Institute, University of Southern California, Marina del Rey, CA. 90292, March 1984.




Figure 1. Geometry of free space optical interconnection using a holographic optical element (HOE).




Figure 3. Schematic diagram of an optical input gate circuit.






Figure 2. Cross section of three detector types compatible with the MOSIS fabrication processes.

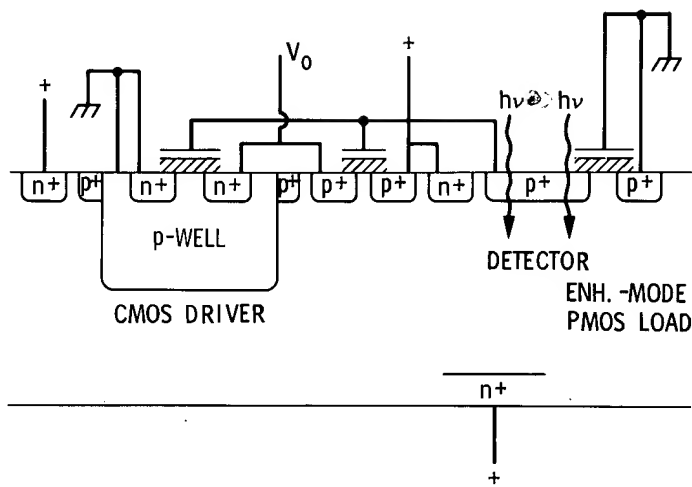


Figure 4. Cross section of an optical gate circuit using a vertical p^+n^- detector.

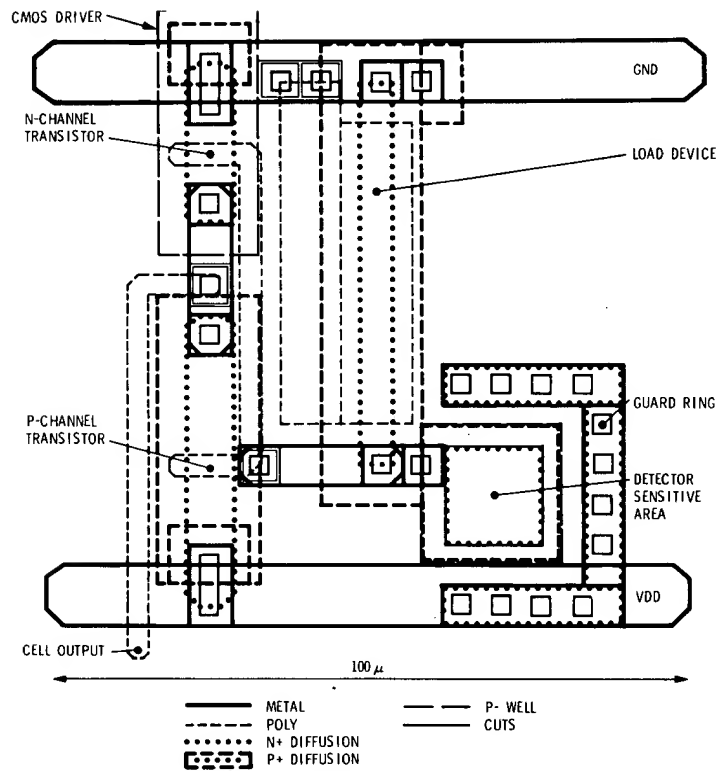


Figure 5. Layout of a CMOS optical input gate incorporating a vertical photodiode detector.




Figure 6. Test chip layout to be used for evaluation of HOE and optical gates.